arXiv:1702.08160v1 [cs.CV] 27 Feb 2017

HashBox: Hash Hierarchical Segmentation exploiting Bounding Box Object Detection

Joachim Curto*,1,3, Irene Zarza*,1,3, Alexander J. Smola[2] and Luc Van Gool[1]
[1]Computer Vision Lab, ETH Zurich, Switzerland
[2]Machine Learning Department, Carnegie Mellon University, PA 15213, USA
[3]Department of Electronic Engineering, City University of Hong Kong, Hong Kong SAR, China

{curto,zarza,vangool}@vision.ee.ethz.ch, alex@smola.org

*Both authors contributed equally.

Abstract

We propose a novel approach to address the Simultaneous Detection and Segmentation problem introduced in [8]. Using the hierarchical structures first presented in [1] we use an efficient and accurate procedure that exploits the hierarchy feature information using Locality Sensitive Hashing. We build on recent work that utilizes convolutional neural networks to detect bounding boxes in an image (Faster R-CNN [11]) and then use the top similar hierarchical region that best fits each bounding box after hashing, we call this approach HashBox. We then refine our final segmentation results by automatic hierarchy pruning. HashBox introduces a train-free alternative to Hypercolumns [7]. We conduct extensive experiments on Pascal VOC 2012 segmentation dataset, showing that HashBox gives competitive state-of-the-art object segmentations.

1. Introduction

Detection and Segmentation are key components in any computer vision toolbox. In this paper we present a hashing technique to segment an object given its bounding box and therefore attain both Detection and Segmentation simultaneously. At its heart lies a novel way to retrieve and generate a high quality segmentation, which is crucial for a wide variety of computer vision applications. Simply put, we use a state-of-the-art convolutional network to detect the objects, but a hashing technique build on top a high quality hierarchy of regions to generate the segmentations.

Object Detection and Segmentation are two popular problems in Computer Vision, historically treated as separated tasks. We consider these strongly related vision tasks as a unique one: detecting each object in an image and assigning to each pixel a binary label inside the corresponding bounding box.

We address the problem with a surprising different technique that deviates from the current norm of using proposal object candidates [2]. In semantic segmentation the need for rich information models that entangle some kind of notion from the different parts that constitute an object is exacerbated. To alleviate this issue we build on the use of a the hierarchical model in [1] and explore the rich space of information of the Ultrametric Contour Map in order to find the best possible semantic segmentation of the given object. For this task, we exploit bounding boxes to facilitate the search task. Hence, we simply hash the bounding box patches and retrieve closest nearest neighbors to the given objects, this is where the name HashBox comes. Using this simple but effective technique we get the segmentation mask which is then refined using Hierarchical Section Pruning.

We start from a bounding box detector and refine the object support, as in Arbeláez *et al*. in Hypercolumns [7]. We propose here a train-free similarity hashing alternative to their approach.

We present a simple yet effective module that leverages the need for a training step and can provide segmentations after any given detector. Our approach is first to use an state-of-the-art region based CNN detector, Faster-RCNN, as prior step to guide the segmentation process. Then we compute the upper bound results when the bounding box detection is perfect, i.e using the ground truth bounding boxes.

Outline: We begin next with a high-level description of the proposed method and develop further the idea to

propose Hierarchical Section Hashing and Hierarchical Section Pruning in Section 1 and Section 3. Current related work follows in Section 2 and we conclude with the evaluation metrics in Section 4.

We start with a primer. HashBox consists on the following main blocks:

- **Bounding Box Object Detection** We use a convolutional neural network system [11] to detect all the objects in an image and generate the corresponding bounding boxes. We consider a detected object in an image as each output candidate thresholded by the class level score (benchmarks specifications in section 4).
- **Hierarchical Image Representation** We represent the image as a hierarchical region tree representation based on the Ultrametric Contour Map of Arbeláez *et al*. [1].
- **Similarity Hashing** We develop Hierarchical Section Hashing based on the LSH technique of [3].
- **Region Refinement** We refine the segmentations by the use of Hierarchical Section Pruning.
- **Evaluation** We evaluate the results on the Pascal VOC 2012 Segmentation dataset [4] using the Jaccard index metric, which measures the average best overlap achieved by a segmentation mask for a ground truth object.

This work is inspired on how humans segment images: they first localize the objects they want to segment, they carefully inspect the object on the image by the use of their visual system and finally they choose the region that belongs to the body of that particular object. We believe that although the problem of detection has to be solved by the use of deep learning techniques based on convolutional neural networks, the problem of segmenting those objects is of a different nature and can be best understood by the use of hashing.

Our main contributions are presented as follows:

- Novel approach to solve the segmentation task exploiting bounding box object detection using similarity hashing.
- Use of hierarchical structures which are rich on semantic meaning instead of other current state-of-the-art techniques such as proposal object candidate generation.
- No need of training data for the segmentation task under bounding box detection framework, i.e train-free accurate segmentations.
- State-of-the-art results.

To our knowledge, we are the first to provide a segmentation solution based on a hashing technique. This approach leverages the need to optimize over a high dimensional space.

Despite the success of region proposal methods in detection, they have in turn arisen as the main computational bottleneck of these approaches. Yet unlike the latter, hierarchical structures derived from the UCM are in comparison inexpensive to compute and store. While we continue to use a very fast region-based convolutional neural network (R-CNN) to solve the detection task, we propose to solve the segmentation problem by exploring efficiently the space generated by a hierarchical image representation.

2. Related Work

Some recent works [8] present Object Detection and Segmentation as a single problem. The SDS task requires to detect and segment every instance of a category in the image. Our work is however more closely related to the Hypercolumns approach in [7], where they go from bounding boxes to segmented masks. Our approach is related in the sense that we propose an alternative that does not require a training step and can be used as an off-the-shelf high quality segmenter.

For semantic segmentation, there has been several approaches where they guide the segmentation process by the use of a prior detector. Also our segmenter starts rather than from raw pixels, Long et al [9], or bounding box proposals as in Girshick et al. [6] and Hariharan et al. [8] , from a set of hierarchical regions given by the UCM structure. Some other work rely on a superpixel representation Mostajabi et al. [10]. This is a different approach than other settings that works directly on a different representation.

3. HashBox

We delve into the details of the HashBox construction.

3.1. Bounding Box Object Detection

We begin by using the Faster-RCNN object detector proposed by Ren et al in [11], which is in turn based on Fast R-CNN by Girshick [5]. Faster R-CNN introduces a Region Proposal Network (RPN) for the task of generating detection proposals and then solves the detection task by the use of a Fast-RCNN detector. They train a CNN on ImageNet

Classification and fine-tune the network on the Pascal detection set. For our experiments, we use the network trained on Pascal VOC 2007 and Pascal VOC 2012 and evaluate the results on the Pascal 2012 eval set. We use the very deep VGG-16 model [12].

3.2. Hierarchical Image Representation

We consider the hierarchical image representation described in [2]. Considering a segmentation of an image into regions that partition its domain $S = \{S_i\}_i$. A segmentation hierarchy is a family of partitions $\{S^*, S^1, \ldots, S^L\}$ such that: (1) S^* is the finest set of superpixels, (2) S^L is the complete domain, and (3) regions from coarse levels are unions of regions from fine levels.

3.3. Hierarchical Section Hashing

In this paper we introduce a novel segmentation algorithm that exploits current bounding box information to automatically select the best hierarchical region that segments the image. We introduce Hierarchical Section Hashing (HSH) which is in turn based on Locality-Sensitive Hashing (LSH). This algorithm helps us surpass the problem of computational complexity of the k-nearest neighbor rule and allows us to do a fast approximate neighbor search in the hierarchical structure of [1].

We can summarize HSH as follows

- Detect bounding boxes on an image using an state-of-the-art convolutional neural network structure such as Faster-RCNN [11].
- Construct a hierarchical image map by using the Ultrametric Contour Map (UCM) and expressing the result as a hierarchical region tree.
- Each hierarchical region is indexed by a number of hash tables using LSH and then constructing a HSH map.
- Each bounding box is hashed into the HSH map to retrieve the approximate nearest neighbor in sublinear time.

We can understand that HashBox has two main steps: first we 'train' the HSH map with all the hierarchical regions of the image. Then we 'test' the HSH map with all the detected bounding boxes in the image to retrieve the approximate nearest neighbors that segment each of the objects in the image. The novelty of this approach is that it provides the best hierarchical region provided by the UCM structure that segments the object image. We say HashBox exploits bounding box object detection because it relies on the correct detection of the object of Faster-RCNN detector.

3.4. Hierarchical Section Pruning

The final piece is to refine the segmentations given by the HSH map by using what we call Hierarchical Section Pruning (HSP).

The HSP procedure can be summarized as follows:

- Once a segmentation mask has been selected for all the objects in the given image, and their bounding boxes recomputed, the bounding box overlap ratio for all box pair combinations, according to the intersection over union criterium, is performed.
- Those masks that present overlap with other object masks on the same image are hierarchically unselected. We always proceed to unselect the low level hierarchy regions, which by construction enclose a smaller region area and thus a single segmented object, from the high level hierarchy region, which enclose more than one object and a bigger image area.
- Finally, isolated pixels on the mask are erased to preserve a single connected segmentation.

HSP is based on the fact that each segmentation mask represent a node on the hierarchical region tree constructed from the UCM. Therefore, hierarchical sections containing more than one object represent higher level nodes on the hierarchy. When HSP is applied, low level hierarchical regions are unselected from the high level hierarchical sections and therefore replaced by mid-low level sections on the same region tree structure that represents a single object or a smaller area of the image.

HashBox relies on the prior detection and therefore availability of bounding boxes for all the objects in a given image. The latter can be very useful as can be thought we can use HashBox as a simple and effective technique to provide high quality segmentations of still images after any available bounding box detector. Likewise, you get train-free off-the-shelf accurate segmentations for any given bounding box detection method.

3.5. Locality Sensitive Hashing

Our goal is to retrieve the k-nearest neighbors of a given hierarchy vector, which we call *image code*. In this setup we are limited by the curse of dimensionality and therefore using an exact search is inefficient. Our approach uses an approximate nearest neighbor technique: Locality Sensitive Hashing (LSH).

A Locality Sensitive Hash function maps $x \rightarrow h(x)$ such that the similarity between $(\mathbf{x}, \mathbf{y})$ is preserved as

$$\left| \frac{d(h(\mathbf{x}), h(\mathbf{y}))}{D(\mathbf{x}, \mathbf{y})} - 1 \right| \leq \epsilon \tag{1}$$

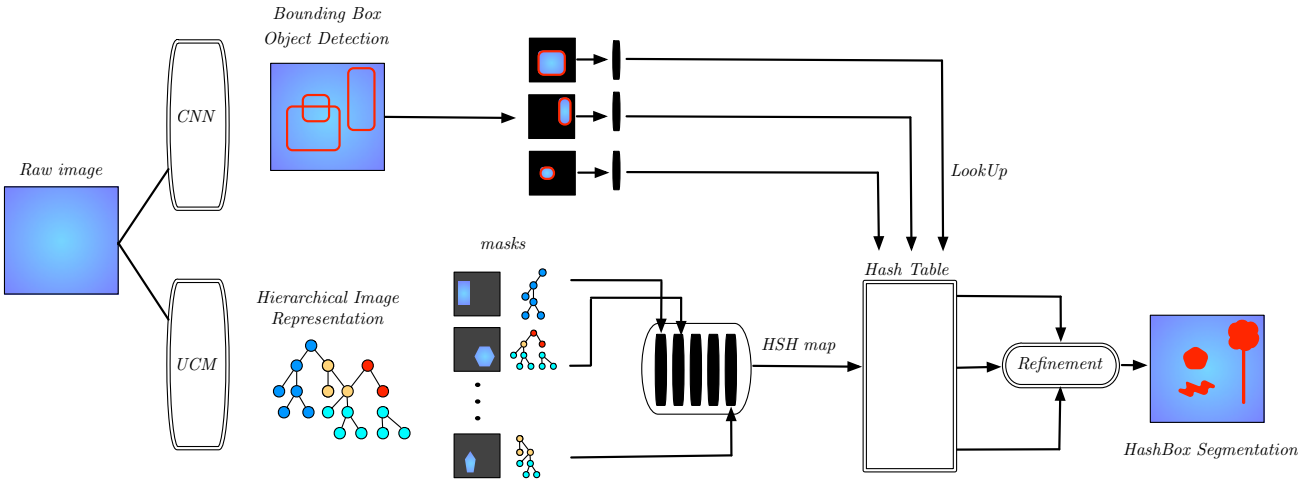

Figure 1. **Hashbox**. We construct a hierarchical image representation based on the UCM and 'train' the HSH map by hashing each of the parent partition node regions. To retrieve a segmentation mask, we 'test' the HSH map by doing a lookup of the detected bounding box region, i.e. a fast approximate nearest neighbor search on the hierarchical structure, and finally refine the result through HSP.

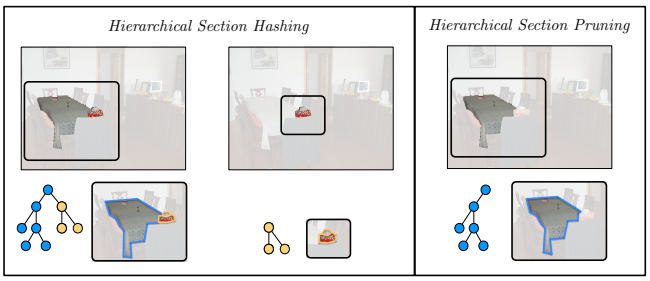

Figure 2. Visual example of Hierarchical Section Hashing and Hierarchical Section Pruning

which is not possible for all $D(\mathbf{x}, \mathbf{y})$ but available for instance for euclidean metrics.

We build on the LSH work of [3]. LSH is a randomized hashing scheme, investigated with the primary goal of $\epsilon - R$ neighbor search. Its main constitutional block is a family of locality sensitive functions. We can define $\mathcal{H}$ of functions $h : \mathcal{X} \rightarrow \{0, 1\}$ is a (p_1, p_2, r, R)-sensitive if, for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$,

$$Pr_{h \sim U[\mathcal{H}]}(h(\mathbf{x}) = h(\mathbf{y}) | |\mathbf{x} - \mathbf{y} \leq r||) \geq p_1, \quad (2)$$

$$Pr_{h \sim U[\mathcal{H}]}(h(\mathbf{x}) = h(\mathbf{y}) | |\mathbf{x} - \mathbf{y} \leq R||) \geq p_2, \quad (3)$$

Where these probabilities are chosen from a random choice of $h \in \mathcal{H}$.

Algorithm 1 gives a simple description of the LSH construction algorithm for the given case when the distance of interest is L_1, which is the one in use in Hashbox. The family $\mathcal{H}$ in this case contains axis-parallel stumps, which means the value of an $h \in \mathcal{H}$ is generated by taking a simple dimension $d \in \{1, \ldots, dim(\mathcal{X})\}$ and thresholding it with some T:

$$h^{LSH} = \begin{cases} 1 & \text{if } x_d \leq T, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

An LSH function $g : \mathcal{X} \rightarrow \{0, 1\}^k$ is formed by independently k function $h_1, \ldots, h_k \in \mathcal{H}$.

That is, we can understand that an example in our hierarchical partition $S_i \in \mathcal{S}$ provides a k-bit hash key

$$g(S_i) = [h_1(S_i), \ldots, h_k(S_i)]. \quad (5)$$

This process is repeated l times and produces l independently constructed hash functions $g_1, \ldots, g_l$. The available reference ('training') data S are indexed by each one of the l hash functions, producing l hash tables, i.e each of the S_i hierarchical partitions generated by all the corresponding parents of the hierarchical tree structure.

Once the LSH data structure has been constructed it can be used to perform a very efficient search for approximate neighbors in the following way.

When a query S_0 arrives, we compute its key for each hash table j, and record the examples $C = \{S_1^l, \ldots, S_{n_l}^l\}$ resulting from the lookup with that key. In other words, we find the 'training' examples that fell in the same bucket of the l-th hash table to which S_0 would fall. These l lookup operations produce a set of candidate matches, $C = \cup_{j=1}^{l} C_j$. If this set is empty, the algorithm reports that and

stops. Otherwise, the distances between candidate matches and S_0 are explicitly evaluated, and the examples that match the search criteria, which means that are closer to S_0 than $(1+\epsilon)R$, are reported.

Algorithm 1 LSH construction

Given: Dataset $X = [\mathbf{x}_1, \mathbf{x}_N], \mathbf{x}_i \in \mathbb{R}^{dim((X))}$.
Given: Number of bits k, number of tables l.
Output: A set of

1: **for** $j = 1, \ldots, l$ **do**
2: **for** $i = 1, \ldots, k$ **do**
3: Randomly (uniformly) draw

$$d \in \{1, \ldots, dim(\mathcal{X})\}.$$

4: Randomly (uniformly) draw

$$\min\{\mathbf{x}_d\} \leq v \leq \max \mathbf{x}_{(d)}.$$

5: Let h_i^j be the function $\mathcal{X} \rightarrow \{0, 1\}$ defined by

$$h_i^j(\mathbf{x}) = \begin{cases} 1 & \text{if } x_{(d)} \leq v, \\ 0 & \text{otherwise.} \end{cases}$$

6: The j-th LSH function is $g_j = [h_1^j, \ldots, h_k^j]$.

4. Evaluation and Results

We extensively evaluate HashBox on PASCAL VOC 2012 validation set.

4.1. Jaccard index metric results

In table 1 we show the results of the Jaccard index metric. This measure represents the average best overlap achieved by a candidate for a ground truth object.

HashBox with 'fast' UCM Jaccard at instance level 45.24 % and Jaccard at class level 43.05 % . Recall at overlap 0.5 is 43.36 %.

References

[1] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *PAMI*, 2011.

[2] P. Arbeláez, J. Pont-Tuset, J. T. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. *CVPR*, 2014.

[3] M. Charikar. Similarity estimation techniques from rounding algorithms. *STOC*, 2002.

[4] M. Everingham, L. V. Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.

[5] R. Girshick. Fast r-cnn. *ICCV*, 2015.

[6] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *CVPR*, 2014.

[7] B. Hariharan, P. Arbeláez, R. Girshick, and J.Malik. Hypercolumns for object segmentation and fine-grained localization. *CVPR*, 2015.

[8] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Simultaneous detection and segmentation. *ECCV*, 2014.

[9] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *CVPR*, 2015.

[10] M. Mostajabi, P. Yadollahpour, and G. Shakhnarovich. Feedforward semantic segmentation with zoom-out features. *CVPR*, 2015.

[11] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*, 2015.

[12] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.

Figure 3. **Top detections**. Top: ground truth segmentations from Pascal VOC 2012. Bottom: HashBox segmentations.

	Aeroplane	Bicycle	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	MBike	Person	Plant	Sheep	Sofa	Train	TV	Global
HashBox 'fast'	45.4	27.5	55.9	44.2	42.0	43.2	41.3	66.3	31.4	57.2	42.3	63.3	43.8	43.6	40.9	40.6	57.2	51.2	48.0	54.1	45.2
HashBox 'fast'	33.3	18.5	48.1	37.5	40.7	45.1	39.4	59.9	23.3	51.0	43.3	60.4	39.8	43.1	34.6	37.2	51.0	47.0	53.6	54.2	43.1

Table 1. **VOC12 val set**. Per-class and global Jaccard index at instance level(J_i)